p-Causality: Identifying Spatiotemporal Causal Pathways for Air Pollutants with Urban Big Data

Julie Yixuan Zhu[1,3,*] Chao Zhang[2,3,*] Shi Zhi[2] Victor O.K. Li[1] Jiawei Han[2] Yu Zheng[3,+]
[1]Department of Electrical and Electronic Engineering, the University of Hong Kong, HK
[2]Dept. of Computer Science, University of Illinois at Urbana-Champaign, Urbana, IL, USA
[3]Microsoft Research Asia, Beijing, China
[1]{yxzhu,vli}@eee.hku.hk [2]{czhang82,shizhi2, hanj}@illinois.edu [3]yuzheng@microsoft.com

arXiv:1610.07045v1 [cs.AI] 22 Oct 2016

ABSTRACT

Many countries are suffering from severe air pollution. Understanding how different air pollutants accumulate and propagate is critical to making relevant public policies. In this paper, we use urban big data (air quality data and meteorological data) to identify the *spatiotemporal (ST) causal pathways* for air pollutants. This problem is challenging because: (1) there are numerous noisy and low-pollution periods in the raw air quality data, which may lead to unreliable causality analysis; (2) for large-scale data in the ST space, the computational complexity of constructing a causal structure is very high; and (3) the *ST causal pathways* are complex due to the interactions of multiple pollutants and the influence of environmental factors. Therefore, we present *p-Causality*, a novel pattern-aided causality analysis approach that combines the strengths of *pattern mining* and *Bayesian learning* to efficiently and faithfully identify the *ST causal pathways*. First, *Pattern mining* helps suppress the noise by capturing frequent evolving patterns (FEPs) of each monitoring sensor, and greatly reduce the complexity by selecting the pattern-matched sensors as "causers". Then, *Bayesian learning* carefully encodes the local and ST causal relations with a Gaussian Bayesian network (GBN)-based graphical model, which also integrates environmental influences to minimize biases in the final results. We evaluate our approach with three real-world data sets containing 982 air quality sensors, in three regions of China from 01-Jun-2013 to 19-Dec-2015. Results show that our approach outperforms the traditional causal structure learning methods in time efficiency, inference accuracy and interpretability.

Keywords

Causality; pattern mining, Bayesian learning; spatiotemporal (ST) big data; urban computing.

1. INTRODUCTION

Recent years have witnessed the air pollution problem becoming a severe environmental and societal issue around the world. For example, in 2015, the average concentration of PM2.5 in Beijing is greater than 150, classified as hazardous to human health by the World Health Organization, on more than 46 days. On Dec 7th 2015, the Chinese government issues the first red alert because of the extremely heavy air pollution, leading to suspended schools, closed construction sites, and traffic restrictions. Though many ways have been deployed to reduce the air pollution, the severe air pollution in Beijing has not been significantly alleviated.

Identifying the causalities has become an urgent problem for mitigating the air pollution and suggesting relevant public policy making. Previous research on the air pollution cause identification mostly relies on chemical receptor [15] or dispersion models [14]. However, these approaches often involve domain-specific data collection which is labor-intensive, or require theoretical assumptions that real-world data may not guarantee. Recently, with the increasingly available air quality data collected by versatile sensors deployed in different regions, and pubic meteorological data, it is possible to analyze the causality of air pollution through a data-driven approach.

The goal of our research is to learn the *spatiotemporal (ST) causal pathways* among different pollutants, by mining the dependencies among air pollutants under different environmental influences. Figure 1 shows two example causal pathways for PM10 in Beijing. Let us first consider the pathway in Figure 1(a). When the wind speed is less than 5 m/s, the high concentration of PM10 in Beijing is mainly caused by SO_2 in Zhangjiakou and PM2.5 in Baoding. In contrast, as shown in Figure 1(b), when the wind speed is larger than 5m/s, PM10 in Beijing is mainly due to PM2.5 in Zhangjiakou and NO_2 in Chengde. Based on this example, we can see the *spatiotemporal (ST) causal pathways* should reflect the following two aspects: 1) the *structural dependency*, which indicates the reactions and propagations of multiple pollutants in the ST space; and 2) the *global confounder*, which denotes how different environmental conditions could lead to different causal pathways.

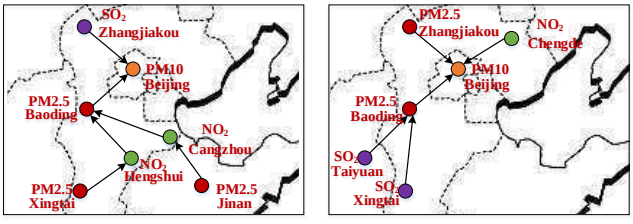

(a) Causal pathways (wind < 5m/s) (b) Causal pathways (wind > 5m/s)

Figure 1: An illustration of identifying causal pathways.

However, identifying the *ST causal pathways* from big air quality and meteorological data is not trivial because of the following challenges. *First*, not all air pollution data are useful for causal-

*Equal contribution. The first authors were interns supervised by the correspondence author in MSRA.
+Correspondence author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.
©20XX ACM X-XXXXX-XX-X/XX/XX.

ity analysis. In the raw sensor-collected air quality data, there are numerous uninteresting fluctuations and noisy variations. Including such data into the causality analysis process is expected to lead to unreliable conclusions. *Second*, the sheer size of the air quality makes the causality analysis difficult. In most air quality monitoring applications, thousands of sensors are deployed at different locations to record the air quality hourly for years. Discovering the ST causal relationships from such a large scale is challenging. *Third*, air pollution causal pathways are complex in nature. The air polluting process typically involves multiple types of pollutants that are mutually interacting, and is subject to local reactions, ST propagations and confounding factors, such as wind and humidity.

Existing data mining techniques for learning the causal pathways can be categorized into two classes: pattern-based [26][20] and Bayesian-based [21][31]. Pattern-based approaches aim to extract frequently occurring phenomena from historical data by applying pattern mining techniques; while Bayesian-based techniques use directed acyclic graphs (DAGs) to encode the causality and then learn the probabilistic dependencies from historical data. Though inspiring results have been obtained by pattern-based and Bayesian-based techniques, both approaches have their merits and downsides. Pattern-based approaches can fast extract a set of patterns (e.g., frequent patterns, contrast patterns) from historical air quality data. Such patterns can capture the intrinsic regularity present in historical air quality data. However, they only provide shallow understanding of the air polluting process, and there are usually a huge number of frequent patterns, which largely limits the usability of the pattern set. On the other hand, Bayesian-based approaches depict the causal dependencies between multiple air pollutants in a principled way. However, the performance of Bayesian-based models is highly dependent on the quality of the training data. When there exist massive noise and data sparsity, as the case of the air quality data, the performance of the Bayesian-based models is limited. Besides, Bayesian-based approaches are limited by high computational cost [4] and the impact of confounding [25].

We propose *p-Causality*, which combines *pattern mining* with *Bayesian learning* to unleash the strengths of both. We claim *p-Causality* is essential for *ST causal pathway* identification, with the contributions listed as below:

• First, a pattern mining module is designed to extract the *frequent evolving patterns* (FEPs) for each sensor, which helps suppress the noise and reduce the computational complexity in the subsequent causal structure learning.

• Second, a Bayesian learning module is proposed to systematically encode and learn the causal pathways for air pollutants in the ST space, based on the high-quality data selected in the previous module. We also integrate the environmental factors to minimize the biases in the final results.

• Third, we have carefully evaluated our proposed approach on three real data sets with 2.5 years' air quality and meteorological data collected from hundreds of cities in China. Our results show that the proposed approach is significantly better than the existing baseline methods in time efficiency, inference accuracy and interpretability.

2. FRAMEWORK

In this section, we first describe the problem of identifying spatio-temporal causal pathways for air pollutants, and then introduce the framework of *p-Causality*.

Let $\mathcal{S} = \{s_1, s_2, \ldots, s_{\mathrm{N}}, \ldots\}$ be the location set of the air quality monitoring sensors deployed in a geographical region. Each sensor is deployed at a location $s_n \in \mathcal{S}$ to periodically measure the target condition around it. All sensors have synchronized measurements over the time domain $\mathcal{T} = \{1, 2, \ldots, \mathrm{T}\}$, where each $t \in \mathcal{T}$ is a timestamp. We also consider a set $\mathcal{C} = \{c_1, c_2, \ldots, c_{\mathrm{M}}\}$ of pollutants. Given $c_m \in \mathcal{C}$, $s_n \in \mathcal{S}$, and $t \in \mathcal{T}$ ($1 \le m \le \mathrm{M}, 1 \le n \le \mathrm{N}, 1 \le t \le \mathrm{T}$), we use $P_{c_m s_n t}$ to denote the measurement of pollutant c_m at location s_n and timestamp t. In addition, we also have the meteorological data at timestamp t for the entire geographical region, denoted as $\boldsymbol{E_t}$, as a vector of environmental factors. Using the air pollutant measurements and meteorological data, we aim to identify faithful causal relationships among different pollutants at different locations.

Figure 2 shows the framework of our proposed approach *p-Causality*. It consists of two main modules: pattern mining and Bayesian Network Learning, detailed as follows.

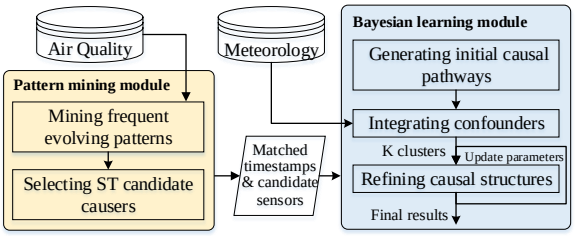

Figure 2: The framework of our approach.

Pattern Mining Module: This module first extracts the *frequent evolving patterns* (FEPs) for each sensor. The FEPs essentially capture the air quality changing behaviors that frequently appear on the target sensor. By mining all FEPs from the historical air quality data, this module efficiently captures the regularity in raw data and largely reduces the noise (Section 3.1 and 3.2). Afterwards, we examine the pattern-based similarities between locations to select candidate causers for each target sensor. By comparing the FEPs occurring on different sensors, we can obtain a shallow understanding of the causal relationships between different sensors, which can be further utilized to simplify learning the causal structures (Section 3.3).

Bayesian Learning Module: By using the matched timestamps of the extracted FEPs at different sensors, together with the selected candidate sensors in the pattern mining module, this module further trains high-quality causal pathways from the large-scale air quality and context data in an effective and scalable way. We first generate the initial causal pathways from the selected candidate causers, taking into account both the local interactions of multiple air pollutants and the ST propagations (Section 4.1). Then to minimize the impact of confounding (Section 4.2), we integrate the confounders (e.g., wind, humidity) into the a Gaussian Bayesian Network (GBN)-based graphical model. Last, we refine the parameters and structures of the Bayesian network to generate the final causal pathways (Section 4.3).

3. THE PATTERN MINING MODULE

3.1 Frequent Evolving Pattern

To capture the frequent evolving behaviors of each sensor, we define *frequent evolving pattern* (FEP), an adaption of the classic sequential pattern concept [22]. As the sequential patterns are defined on transactional sequences, we first discretize the raw air quality data. Given a pollutant c_m at sensor s_n, the measurements of c_m at s_n over the time domain $\mathcal{T}$ form a time series. We discretize the time series as follows: (1) partition it by day to obtain a collection of daily time series, denoted as $P_{c_m s_n}$; and (2) for each

daily time series $\langle (p_1, t_1), (p_2, t_2), \ldots, (p_l, t_l) \rangle$, map every real-value measure p_i ($1 \leq i \leq l$) to a discrete level $\hat{p}_i$ using *symbolic approximation aggregation* [17]. After discretization, we obtain a database of symbolic sequences, as defined in Definition 1.

DEFINITION 1 (SYMBOLIC POLLUTION DATABASE). *For pollutant c_m and sensor s_n, the symbolic pollution database $\hat{P}_{c_m s_n}$ is a collection of daily sequences. Each sequence $d \in \hat{P}_{c_m s_n}$ has the form $\langle (\hat{p}_1, t_1), (\hat{p}_2, t_2), \ldots, (\hat{p}_l, t_l) \rangle$ where an element $(\hat{p}_i, t_i)$ means the pollution level of c_m at sensor s_n and time t_i is $\hat{p}_i$.*

Given the database $\hat{P}_{c_m s_n}$, our goal is to find frequent evolving behaviors of s_n regarding c_m. Below, we introduce the concepts of *evolving sequence* and *occurrence*.

DEFINITION 2 (EVOLVING SEQUENCE). *A length-k evolving sequence T has the form $T = \hat{p}_1 \xrightarrow{\Delta t} \hat{p}_2 \xrightarrow{\Delta t} \cdots \xrightarrow{\Delta t} \hat{p}_k$, where (1) $\forall i > 1, \hat{p}_{i-1} \neq \hat{p}_i$ and (2) Δt is the maximum transition time between consecutive records.*

DEFINITION 3 (OCCURRENCE). *Given a daily sequence $d = \langle (\hat{p}_1, t_1), (\hat{p}_2, t_2), \ldots, (\hat{p}_l, t_l) \rangle$ and an evolving sequence $T = \hat{p}_1 \xrightarrow{\Delta t} \hat{p}_2 \cdots \xrightarrow{\Delta t} \hat{p}_k$ ($k \leq l$), T occurs in d (denoted as $T \sqsubseteq d$) if there exist integers $1 \leq j_1 < j_2 < \cdots < j_k \leq l$ such that: (1) $\forall 1 \leq i \leq k$, $\hat{p}_{j_i} = \hat{p}_i$; and (2) $\forall 1 \leq i \leq k-1$, $0 < t_{j_{i+1}} - t_{j_i} \leq \Delta t$.*

For clarity, we denote an evolving sequence $\hat{p}_1 \xrightarrow{\Delta t} \hat{p}_2 \cdots \xrightarrow{\Delta t} \hat{p}_k$ as $\hat{p}_1 \rightarrow \hat{p}_2 \cdots \rightarrow \hat{p}_k$ when the context is clear. Now, we proceed to define *support* and *frequent evolving pattern*.

DEFINITION 4 (SUPPORT). *Given $\hat{P}_{c_m s_n}$ and an evolving sequence T, the support of T is the number of days that T occurs,* i.e.*, $Sup(T) = |\{o | o \in \hat{P}_{c_m s_n} \wedge T \sqsubseteq o\}|$.*

DEFINITION 5 (FREQUENT EVOLVING PATTERN). *Given a support threshold σ, an evolving sequence T is a frequent evolving pattern in database $\hat{P}_{c_m s_n}$ if $Sup(T) \geq \sigma$.*

3.2 The FEP Mining Algorithm

Now we proceed to discuss how to mine all FEPs in any symbolic pollution database. It is closely related to the classic sequential pattern mining problem. However, recall that there are two constraints in the definition of FEP: (1) the consecutive symbols must be different; and (2) the time gap between consecutive records should be no greater than the temporal constraint Δt. A sequential pattern mining algorithm needs to be tailored to ensure these two constraints are satisfied.

We adapt PrefixSpan [22] as it has proved to be one of the most efficient sequential pattern mining algorithms. The basic idea of PrefixSpan is to use short patterns as the prefix to project the database and progressively grow the short patterns by searching for local frequent items. For a short pattern β, the β-projected database $\mathcal{D}_\beta$ includes the postfix from the sequences that contain β. Local frequent items in $\mathcal{D}_\beta$ are then identified and appended to β to form longer patterns. Such a process is repeated recursively until no more local frequent items exist. One can refer to [22] for more details.

Given a sequence α and a frequent item $\hat{p}$, when creating $\hat{p}$-projected database, the standard PrefixSpan procedure generates one postfix based on the first occurrence of $\hat{p}$ in α. This strategy, unfortunately, can miss FEPs in our problem.

EXAMPLE 1. *Let $\Delta t = 60$ and $\sigma = 3$. In the database shown in Table 1, item $\hat{p}_1$ is frequent. The $\hat{p}_1$-projected database generated by PrefixSpan is:*

(1) $d_1/\hat{p}_1 = \langle (\hat{p}_2, 20), (\hat{p}_3, 30) \rangle$
(2) $d_2/\hat{p}_1 = \langle (\hat{p}_2, 30), (\hat{p}_1, 360), (\hat{p}_2, 400), (\hat{p}_3, 420) \rangle$
(3) $d_4/\hat{p}_1 = \langle (\hat{p}_1, 120), (\hat{p}_3, 140), (\hat{p}_2, 150), (\hat{p}_3, 180) \rangle$

The elements satisfying $t \leq 60$ are $(\hat{p}_2, 20)$, $(\hat{p}_3, 30)$ and $(\hat{p}_2, 30)$. No local item is frequent, hence $\hat{p}_1$ cannot be grown any more.

Table 1: An example symbolic pollution database.

Day	Daily sequence
d_1	$\langle (\hat{p}_2, 0), (\hat{p}_1, 10), (\hat{p}_2, 30), (\hat{p}_3, 40) \rangle$
d_2	$\langle (\hat{p}_1, 0), (\hat{p}_2, 30), (\hat{p}_1, 360), (\hat{p}_2, 400), (\hat{p}_3, 420) \rangle$
d_3	$\langle (\hat{p}_2, 0), (\hat{p}_3, 30) \rangle$
d_4	$\langle (\hat{p}_1, 0), (\hat{p}_1, 120), (\hat{p}_3, 140), (\hat{p}_2, 150), (\hat{p}_3, 180) \rangle$
d_5	$\langle (\hat{p}_2, 50), (\hat{p}_2, 80), (\hat{p}_3, 120), (\hat{p}_1, 210) \rangle$

To overcome this, given a sequence α and a frequent item $\hat{p}$, we generate a postfix for every occurrence of $\hat{p}$.

EXAMPLE 2. *Also for Example 1, if we generate a postfix for every occurrence of $\hat{p}_1$, the $\hat{p}_1$-projected database is:*

(1) $d_1/\hat{p}_1 = \langle (\hat{p}_2, 20), (\hat{p}_3, 30) \rangle$
(2) $d_2/\hat{p}_1 = \langle (\hat{p}_2, 30), (\hat{p}_1, 360), (\hat{p}_2, 400), (\hat{p}_3, 420) \rangle$
(3) $d_2/\hat{p}_1 = \langle (\hat{p}_2, 40), (\hat{p}_3, 60) \rangle$
(4) $d_4/\hat{p}_1 = \langle (\hat{p}_1, 120), (\hat{p}_3, 140), (\hat{p}_2, 150), (\hat{p}_3, 180) \rangle$
(5) $d_4/\hat{p}_1 = \langle (\hat{p}_3, 20), (\hat{p}_2, 30), (\hat{p}_3, 60) \rangle$

The items $\hat{p}_2$ and $\hat{p}_3$ are frequent and meanwhile satisfy the temporal constraint, thus longer patterns $\hat{p}_1 \xrightarrow{60} \hat{p}_2$ and $\hat{p}_1 \xrightarrow{60} \hat{p}_3$ are found in the projected database.

Using the above projection principle, the projected database includes all postfixes to avoid missing patterns under the time constraint. Algorithm 1 sketches our algorithm for mining FEPs. The procedure is similar to the standard PrefixSpan algorithm in [22], except that the aforementioned full projection principle is adopted, and the time constraint Δt is checked when searching for local frequent items.

Algorithm 1: Mining frequent evolving patterns.

```
Input: support threshold σ, temporal constraint Δt, symbolic
       pollution database P̂
Procedure InitialProjection(P̂, σ, Δt)
    ← frequent items in 𝒟;
    foreach item i in do
        S ← φ;
        foreach sequence o in P̂ do
            R ← postfixes for all occurrences of i in o;
            S ← S ∪ R;
        PrefixSpan(i, i, 1, S, Δt);
Function PrefixSpan(α, i_prev, l, S|_α, Δt)
    ← frequent items in S|_α meeting time constraint Δt;
    foreach item i in do
        if i ≠ i_prev then
            α' ← append i to α;
            Build S|_α' using full projection;
            Output α';
            PrefixSpan(α', i, l + 1, S|_α', Δt);
```

The output of Algorithm 1 is the set of all FEPs for the given database, along with the occurring timestamps for each FEP. As an

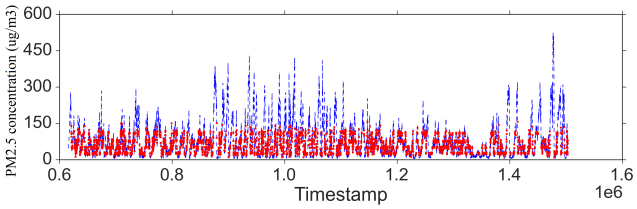

Figure 3: An illustration of the pattern-matched timestamps. The blue dashed lines represents the PM2.5 time series in Beijing during a two-year period, and the red points denote the timestamps at which a certain FEP has occurred ($\delta = 0.1$).

example, Figure 3 shows the raw PM2.5 time series in Beijing during a two-year period. After mining FEPs on the symbolic pollution database, we mark the timestamps at which the FEPs occur. One can observe that, the FEPs can effectively capture the regularly appearing evolvements of PM2.5 in Beijing. Because of the support threshold and the evolving constraint, infrequent sudden changes and uninteresting fluctuations are all suppressed.

3.3 Finding Candidate Causers

After discovering the FEPs, next step is leverage them to extract the candidate causers for each sensor. Consider two sensors s and s', let us use $TS(s)$ and $TS(s')$ to denote the sets of pattern starting timestamps for s and s', respectively. Below, we introduce the *pattern match* relationship.

DEFINITION 6 (PATTERN MATCH). *Let $t_{s'} \in TS(s')$ be a timestamp at which a pattern happens on s'. For a pattern starting timestamp $t_s \in TS(s)$, we say $t_{s'}$ matches t_s if $0 \leq t_s - t_{s'} \leq L$, where L is a pre-specified time lag threshold.*

Informally, the pattern match relation states that when there is a pattern occurring on s', then within some time interval, there is another pattern happening on s. Naturally, if s' has a strong causal effect on s, then most timestamps in $TS_{s'}$ will be matched by TS_s, and vice versa. Based on TS_s and $TS_{s'}$, we proceed to introduce *match precision* and *match recall* to quantify the correlation between s and s'.

DEFINITION 7 (MATCH PRECISION). *Given TS_s and $TS_{s'}$, we define the matched timestamp set of $TS_{s'}$ as $M_{s'} = \{t_{s'} | t_{s'} \in TS_{s'} \wedge \exists t_s \in TS_s, match(t_s, t_{s'}) = True\}$. With $M_{s'}$ and $TS_{s'}$, we define the precision of s' matching s as:*

$$P(s, s') = |M_{s'}|/|TS_{s'}|$$

DEFINITION 8 (MATCH RECALL). *Given TS_s and $TS_{s'}$, we define the matched timestamp set of TS_s as $M_s = \{t_s | t_s \in TS_s \wedge \exists t_{s'} \in TS_{s'}, match(t_s, t_{s'}) = True\}$. With M_s and TS_s, we define the recall of s' matching s as:*

$$R(s, s') = |M_s|/|TS_s|$$

Relying on the concepts of *match precision* and *match recall*, we compute the pattern-based correlation between s and s' as the geometric mean of precision and recall, namely

$$Corr(s, s') = \frac{2 \times P(s, s')}{P(s, s') + R(s, s')}.$$

Now we are ready to describe the process of finding candidate causers for each sensor. Given the set of all sensors and their pattern-starting timestamps, our goal is to find the candidate causers for each sensor. Consider a target sensor s, we say another sensor s' is a candidate causer for s if s' satisfies two constraints: (1) the distance between s and s' is no larger than a distance threshold δ_g; and (2) the pattern correlation between s and s' is no less than a correlation threshold δ_p. Given the pattern-starting timestamps that are ordered chronologically, the retrieval of the candidate causers can be easily done by sequentially scanning the two timestamp lists to find pattern-matched pairs.

Figure 4 illustrates eight examples of selected candidate causers. For PM2.5 in Beijing, we reduce the number of candidate sensors to $X = 4 \sim 7$ from overall $|\mathcal{S}| = 61$ sensors in North China. Note that China is a country with monsoon climate, the candidate sensors show quite similar geo-locations in four seasons. We therefore separate the training data into four groups based on seasons, to better diagnose causalities for the air pollutants in China.

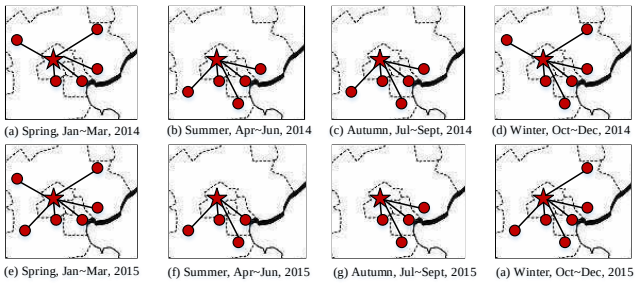
(a) Spring, Jan~Mar, 2014
(b) Summer, Apr~Jun, 2014
(c) Autumn, Jul~Sept, 2014
(d) Winter, Oct~Dec, 2014
(e) Spring, Jan~Mar, 2015
(f) Summer, Apr~Jun, 2015
(g) Autumn, Jul~Sept, 2015
(a) Winter, Oct~Dec, 2015

Figure 4: Candidate sensors for Beijing PM2.5 in four seasons. Star: PM2.5 in Beijing. Circles: pollutants at candidate sensors.

4. THE BAYESIAN LEARNING MODULE

4.1 Generating Initial Causal Pathways

To diagnose the complex reactions and dispersions of different air pollutants, both the local and ST dependencies need to be modelled in a fair way. Based on the pattern-matched data and the candidate causers extracted by the pattern mining module, this subsection first introduces the representation of causal pathways in the ST space, and then elaborates how to generate initial causal pathways.

DEFINITION 9 (GAUSSIAN BAYESIAN NETWORK (GBN)). *GBN is a special form of Bayesian network for probabilistic inference with continuous Gaussian variables in a DAG, in which each variable is assumed as linear function of its parents [9].*

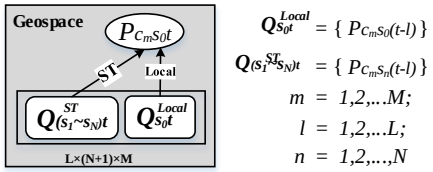

(a) Local and ST dependencies in a GBN (b) Notations

Figure 5: GBN-based causal pathway representation and its notations.

As shown in Figure 5, the ST causal relations of air pollutants are encoded in a GBN-based graphical model, to represent both local and ST dependencies. Here we choose GBN to model the causalities because: 1) GBN provides a simple way to represent the dependencies among multiple pollutants variables, both locally and in the ST space. 2) GBN models continuous variables, while

general Bayesian network (BN) models discrete variables. Since quantifying the continuous urban data into discrete values will introduce errors and more uncertainties for causality analysis, we use GBN instead of BN in our model. 3) The characteristics of urban data fit the GBN model well, since the distribution of 1-hour difference (current value minus the value 1-hour ago) of air pollutants and meteorological data obey Gaussian distribution (verified by $D'Agostino-Pearson$ test [7][30]). In the following sections, normalized 1-hour differences of time series data will be used as inputs for the model.

Specifically, for the target pollutant c_m at sensor s_0-th sensor and timestamp t, denoted as $P_{c_m s_0 t}, m \in [1, \mathrm{M}]$, we capture the dependencies from both the local causal pollutants $\boldsymbol{Q}^{\boldsymbol{Local}}_{\boldsymbol{s_0 t}}$ and the ST causal pollutants $\boldsymbol{Q}^{\boldsymbol{ST}}_{(\boldsymbol{s_1 \sim s_N})\boldsymbol{t}}$. Here $\boldsymbol{Q}^{\boldsymbol{ST}}_{(\boldsymbol{s_1 \sim s_N})\boldsymbol{t}}$ refer to a $1 \times \mathrm{MNL}$ vector of M types of pollutants at N neighborhood sensors $s_1 \sim s_N$ and previous L timestamps that most probably cause the target pollutant in the ST space, i.e. $\boldsymbol{Q}^{\boldsymbol{ST}}_{(\boldsymbol{s_1 \sim s_N})\boldsymbol{t}} = \{P_{c_m s_n (t-l)}\}, m = 1, \ldots, \mathrm{M}; n = 1, \ldots, \mathrm{N}; l = 1, \ldots, \mathrm{L}$. Similarly, $\boldsymbol{Q}^{\boldsymbol{Local}}_{\boldsymbol{s_0 t}}$ is a $1 \times \mathrm{ML}$ vector of pollutants locally at s_0. For example, when we set $\mathrm{L} = 2, \mathrm{M} = 6$, $\boldsymbol{Q}^{\boldsymbol{Local}}_{\boldsymbol{s_0 t}}$ may take values of 12 normalized 1-hour difference time series data, i.e. $\boldsymbol{Q}^{\boldsymbol{Local}}_{\boldsymbol{s_0 t}} = (2, -0.5, 0.8, 0.3, 1, -2, 2.2, 1, 1, 0, -0.5, 0.2)$.

The parents of $P_{c_m s_0 t}$ are denoted as $\boldsymbol{PA}(P_{c_m s_0 t}) = \boldsymbol{Q}^{\boldsymbol{Local}}_{\boldsymbol{s_0 t}} \oplus \boldsymbol{Q}^{\boldsymbol{ST}}_{(\boldsymbol{s_1 \sim s_N})\boldsymbol{t}}$, where $\oplus$ denotes the concatenation operator for two vectors. Based on the definition of GBN, the distribution of $P_{c_m s_0 t}$ conditioned on $\boldsymbol{PA}(P_{c_m s_0 t})$ obeys Gaussian distribution:

$$Pr(P_{c_m s_0 t} = p_{c_m s_0 t} | \boldsymbol{PA}(P_{c_m s_0 t})) \sim \mathcal{N}(\mu_{c_m s_0 t} + \Sigma_{n=0}^{\mathrm{N}} \Sigma_{l=1}^{\mathrm{L}} \Sigma_{m=1}^{\mathrm{M}} a_{m(\mathrm{N}\mathrm{L}+l)} (p_{c_m s_n (t-l)} - \mu_{c_m s_n (t-l)}), \Sigma(\epsilon_{c_m s_0 t})) \quad (1)$$

$\mu_{c_m s_0 t}$ is the marginal mean for $P_{c_m s_0 t}$. Σ denotes the covariance operator. $\boldsymbol{A} = \{a_m(nL + l)\}, (m = 1, \ldots, \mathrm{M}; n = 0, 1, \ldots, \mathrm{N}; l = 1, \ldots, \mathrm{L})$ is the coefficient for the linear regression in GBN [9]:

To minimize the uncertainty of $P_{c_m s_0 t}$ given its parents, we need to find N sensors $s_1 \sim s_N$ from the ST space and the parameters $\boldsymbol{A}$ that minimizes the error:

$$\Sigma(\epsilon_{c_m s_0 t}) = \Sigma(P_{c_m s_0 t}) - \boldsymbol{A} \Sigma(\boldsymbol{PA}(P_{c_m s_0 t}))^{-1} \boldsymbol{A}^T \quad (2)$$

Generating the initial causal pathways requires locating N most influential sensors from overall $|\mathcal{S}|$ sensors requires up to $\binom{|\mathcal{S}|}{N}$ trials. Yet given the candidate sensors selected by Section 3.3, we manage to search from a subset ($X \leq |\mathcal{S}|$) sensors with time efficiency and scalability. We further propose a a Granger causality score GC_{score} to generate initial causal pathways, which is defined as:

$$GC_{score}(m, s_0, s_n) = max_{m' \in [1,\mathrm{M}]} \{ |match(t_{(c_m, s_0)}, t_{(c_{m'}, s_n)})| \cdot \frac{|\Sigma(\epsilon_{c_m s_0 t})_1| - |\Sigma(\epsilon_{c_m s_0 t})_2|}{|\Sigma(\epsilon_{c_m s_0 t})_2| \chi^2_{\mathrm{L}}(0.05)} \} \quad (3)$$

where GC_{score} is a χ^2-test score [12] for the predictive causality, with higher score indicating more probable "Granger" causes from M pollutants at sensor s_n to the target pollutant c_m at sensor s_0 [10] ($GC_{score} \leq 1$ means none causality). For variables obeying Gaussian distribution, Granger causality is in the same form as conditional mutual information [2], which has been used successfully for constructing structures for Bayesian networks. Here $|match(t_{(c_m, s_0)}, t_{(c'_m, s_n)})|$ is the number of matched timestamps of FEPs between two time series (pollutant $c_{m'}$ at sensor s_n and pollutant c_m at sensor s_0, see Section 3.3). And $\Sigma(\epsilon_{c_m s_0 t})_1$ and $\Sigma(\epsilon_{c_m s_0 t})_2$ correspond to the variances of the target pollutant $P_{c_m s_0 t}$ conditioned on $\boldsymbol{Q}^{\boldsymbol{Local}}_{\boldsymbol{s_0 t}}$ and $\boldsymbol{Q}^{\boldsymbol{Local}}_{\boldsymbol{s_0 t}} \oplus \boldsymbol{Q}^{\boldsymbol{ST}}_{\boldsymbol{s_n t}}$.

4.2 Integrating Confounders

Recall the example in Figure 1. A target pollutant is likely to have several different causal pathways under different environmental conditions, which indicates the environmental factors (humidity, wind, etc.) can be treated as confounders to the causalities among pollutants. This subsection is designed for integrating the environmental factors into the GBN-based graphical model, to minimize the biases in causality analysis and guarantee the causal pathways are faithful for the government's decision making. We first introduce the definition of confounder and then elaborate the integration.

DEFINITION 10 (CONFOUNDER). *A confounder is defined as a third variable that simultaneously correlates with the cause and effect, e.g. gender K may affect the effect of recovery P given a medicine Q, as shown in Figure 6(a). Ignoring the confounders will lead to biased causality analysis. To guarantee an unbiased causal inference, the cause-and-effect is usually adjusted by averaging all the sub-classification cases of K [21], i.e. $Pr(P|do(Q)) = \Sigma_{k=1}^{K} Pr(P|Q,k)Pr(k)$.*

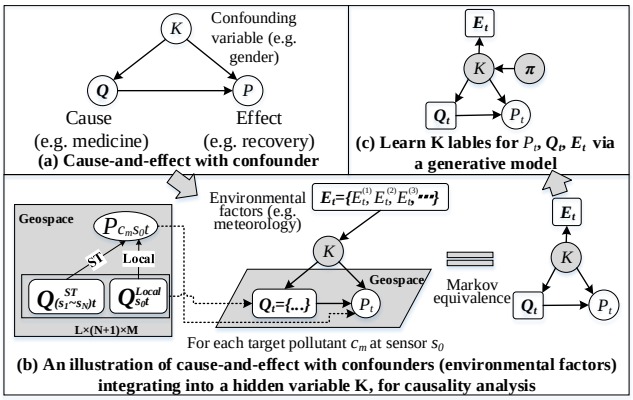

Figure 6: The GBN-based graphical model, integrating confounders to the causal pathway, and converting the model into a generative model

For integrating environmental factors, denoted as $\boldsymbol{E_t} = \{E_t^{(1)}, E_t^{(2)}, \ldots\}$, into the GBN-based causal pathways, one challenge is there can be too many sub-classifications of environmental statuses. For example, if there are 5 environmental factors and each factor has 4 statuses, there will exist $4^5 = 1024$ causal pathways for each sub-classification case. Thus directly integrating $\boldsymbol{E_t}$ as confounders to the cause and effect will result in unreliable causality analysis due to very few sample data conditioned on each sub-classification case. Therefore, we introduce a discrete hidden confounding variable K, which determines the probabilities of different causal pathways from $\boldsymbol{Q_t}$ to $\boldsymbol{P_t}$, as shown in Figure 6(b). The environmental factors $\boldsymbol{E_t}$ are further integrated into K, where $K = 1, 2, ...\mathrm{K}$. In this ways, the large number of sub-classification cases of confounders will be greatly reduced to a small number K, indicating K clusters of the environmental factors $\boldsymbol{E_t}$.

Based on Markov equivalence (DAGs which share the same joint probability distribution [8]), we can reverse the arrow $\boldsymbol{E_t} \rightarrow K$ to $K \rightarrow \boldsymbol{E_t}$, as shown in the right part of Figure 6(b). K determines the distributions of $P, \boldsymbol{Q_t}, \boldsymbol{E_t}$, thus enabling us to learn the distribution of the graphical model from a generative process. To help us learn the hidden variable K, the generative process further introduces a hyper-parameter $\boldsymbol{\pi}$ (as shown in Figure 6(c)) that

determines the distribution of K. Thus the graphical model can be understood as a mixture model under K clusters. We learn the parameters of the graphical model by maximizing the new log likelihood:

$$LL^{gen} = \Sigma_t \Sigma_{k=1}^{\mathrm{K}} ln(Pr(p_t|\boldsymbol{q_t}, k) Pr(\boldsymbol{q_t}|k) Pr(\boldsymbol{e_t}|k) Pr(k|\boldsymbol{\pi})) \tag{4}$$

In determining the number of the hidden variable K, we do not consider too large a K since that will induce much complexity for causality analysis. Also too small a K may not characterize the information contained in the confounders (i.e. meteorology). We observe the 2-D PCA projections of meteorological data (as shown in Figure 7). In three regions, five clusters can characterize the data sufficiently well. Thus we choose $\mathrm{K} = 3 \sim 7$ for learning.

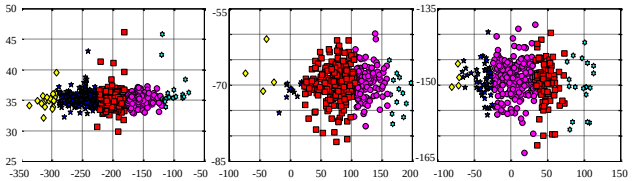

(a) North China (NC) (b) Yangtze River Delta (YRD) (c) Pearl River Delta (PRD)

Figure 7: 2-D PCA projections of 5 clusters of meteorological data in NC, YRD and PRD. The original meteorological data contains five types, i.e. temperature (T), pressure (P), humidity (U), wind speed (WS), and wind direction (WD), with each region divided into 9 grids, thus 45-dimensional.

4.3 Refining Causal Structures

This subsection tries to refine the causal structures and obtain the final causal structures under K clusters. The refining process includes two phases in each iteration: 1) an EM learning (EML) phase to infer the parameters of the model, and 2) a structure reconstruction (SR) phase to re-select the top N neighborhood sensors based on the new learnt parameters and GC_{score}, as illustrated in Algorithm 2.

EML (line 6-18) is an approximation method to learn the parameters $\boldsymbol{\pi}, \boldsymbol{\gamma}, \boldsymbol{A_k}, \boldsymbol{B_k}$ of the graphical model, by maximizing the log likelihood (Equation 4) of the observed data sets via an E-step and a M-step. Here $\boldsymbol{\pi}$ contains the hyper parameters which determine the distribution of K ($\mathrm{T} \times \mathrm{K}$-dimensional). $\boldsymbol{\gamma}$ are posterior probabilities for each monitoring record ($\mathrm{T} \times \mathrm{K}$-dimensional). $\boldsymbol{A_k}, \boldsymbol{B_k}$ are parameters for measuring the dependencies among pollutants and meteorology (K-dimensional). Note that $\boldsymbol{A_k}, \boldsymbol{B_k}$ come in different formats. $\boldsymbol{A_k}$ is the regression parameter for:

$$P_{c_m s_0 t} = \mu_0 + (\boldsymbol{Q^{Local}_{s_0 t}} \oplus \boldsymbol{Q^{ST}_{(s_1 \sim s_N) t}}) \boldsymbol{A_k} + \epsilon_{c_m s_0 t} \tag{5}$$

and $\boldsymbol{B_k} = (\boldsymbol{\mu_{B_k}}, \boldsymbol{\Sigma_{B_k}})$ includes the parameters for the multivariate Gaussian distribution of environmental factors $\boldsymbol{E_t}$. In the E-step, we calculate the expectation of log likelihood (Equation 6) with the current parameters, and the M-step re-computes the parameters.

4.3.1 *E-step*

Given the parameters $\boldsymbol{\pi}$, K, N, $\boldsymbol{A_k}, \boldsymbol{B_k}$, EM assumes the membership probability of $p_t, \boldsymbol{q_t}, \boldsymbol{e_t}$ belonging to the k-th cluster as:

$$\begin{aligned} \gamma_{tk} &= Pr(k|p_t, \boldsymbol{q_t}, \boldsymbol{e_t}) = \frac{Pr(k) Pr(p_t, \boldsymbol{q_t}, \boldsymbol{e_t}|k)}{Pr(p_t, \boldsymbol{q_t}, \boldsymbol{e_t})} \\ &= \frac{\pi_{tk} \mathcal{N}(p_t|\boldsymbol{q_t}, \boldsymbol{A_k}) \mathcal{N}(\boldsymbol{e_t}|\boldsymbol{B_k})}{\Sigma_{j=1}^{K} \pi_{tj} \mathcal{N}(p_t|\boldsymbol{q_t}, \boldsymbol{A_j}) \mathcal{N}(\boldsymbol{e_t}|\boldsymbol{B_j})} \end{aligned} \tag{6}$$

4.3.2 *M-step*

The membership probability in E-step can be used to calculate new parameter values $\boldsymbol{\pi^{new}}, \boldsymbol{A_k^{new}}, \boldsymbol{B_k^{new}}$. We first determine the most likely assignment tag of timestamp t to cluster k, i.e.

$$Tag_t = max_{k \in [1,\mathrm{K}]} \pi_{tk} \tag{7}$$

By integrating the timestamps belonging to each cluster k, we can update $\boldsymbol{A_k^{new}}$ by Equation 5. Then we update $\boldsymbol{B_k}$ by:

$$\begin{aligned} \boldsymbol{\mu^{new}_{B_k}} &= \frac{1}{T_k} \Sigma_{t=1}^{T} \gamma_{tk} \boldsymbol{e_t}, T_k = \Sigma_{t=1}^{T} \gamma_{tk} \\ \boldsymbol{\Sigma^{new}_{B_k}} &= \frac{1}{T_k} \Sigma_{t=1}^{T} \gamma_{tk} (\boldsymbol{e_t} - \boldsymbol{\mu^{new}_{B_k}})(\boldsymbol{e_t} - \boldsymbol{\mu^{new}_{B_k}})^T \end{aligned} \tag{8}$$

In addition, we update π_{tk}^{new} by:

$$\pi_{tk}^{new} = \Sigma_{t=1}^{T} \frac{\gamma_{tk}}{T_k} \tag{9}$$

SR phase (line 19-24) utilizes the parameters provided by the EM learning phase, and re-select the top N neighborhood sensors based on the newly generated GC_{score} for each cluster k. We present a training example (as shown in Figure 8(a)) of learning the causal pathways for Beijing PM2.5 during Jan$\sim$Mar. After 13 training iterations of the EM learning phase and structure reconstruction, we finally obtain $\mathrm{K} = 5$ causal structures under each cluster, with the log likelihood shown in Figure 8(b). For the last iteration, we calculate the percentage of labeled timestamps belonging to each cluster k.

Algorithm 2: Refining the causal structures for each target pollutant c_m at location s_0.

Input: T, K, N, and raining data sets p_t, $\boldsymbol{q_t}$, $\boldsymbol{e_t}$, $t \in [1, \mathrm{T}]$
Output: Refined causal structures for K clusters
1 Initial neighborhood sensors $s_1 \sim s_N$ based on top N GC_{score};
2 **repeat**
3 EML(P_t, $\boldsymbol{Q_t}$, $\boldsymbol{E_t}$, $s_1 \sim s_N$, K) $\rightarrow$ $Log\_likelihood$, $\pi_{tk}, \gamma_{tk}, \boldsymbol{A_k}, \boldsymbol{B_k}$;
4 SR($\boldsymbol{A_k}$, $s_1 \sim s_N$, K)$\rightarrow$ $s'_1 \sim s'_N$, Q';
5 **until** *Log_likeoihood converges*;
6 **Function** *EM_Learning(EML)*(P_t, $\boldsymbol{Q_t}$, $\boldsymbol{E_t}$, $s_1 \sim s_N$, K)
7 **repeat**
8 InitialAssign: K clusters via K-means($\boldsymbol{E_t}$)
9 **foreach** *item* $t = 1$ *to* T **do**
10 **foreach** *item* $k = 1$ *to* K **do**
11 Update π_{tk} by Equation (9);
12 **foreach** *item* $k = 1$ *to* K **do**
13 Update $\boldsymbol{A_k}, \boldsymbol{B_k}$ by Equation (5),(8);
14 **foreach** *item* $t = 1$ *to* T **do**
15 **foreach** *item* $k = 1$ *to* K **do**
16 Update γ_{tk} by Equation (6);
17 **until** *Log likelihood converges*;
18 return: $Log\_likelihood$ and $\pi_{tk}, \gamma_{tk}, \boldsymbol{A_k}, \boldsymbol{B_k}$;
19 **Function** *Structure_Reconstruction(SR)*($\boldsymbol{A_k}$, $s_1 \sim s_N$, K)
20 **foreach** *item* s_n *in All candidate sensors* **do**
21 Compute $GC_{score}(m, s_0, s_n)$ for $s_1 \sim s_N$;
22 Rank GC_{score} and re-select the top N neighborhood sensors $s'_1 \sim s'_N$;
23 Update $Q \rightarrow Q'$ corresponding to $s'_1 \sim s'_N$;
24 return: $s'_1 \sim s'_N$, Q';

5. EXPERIMENTS

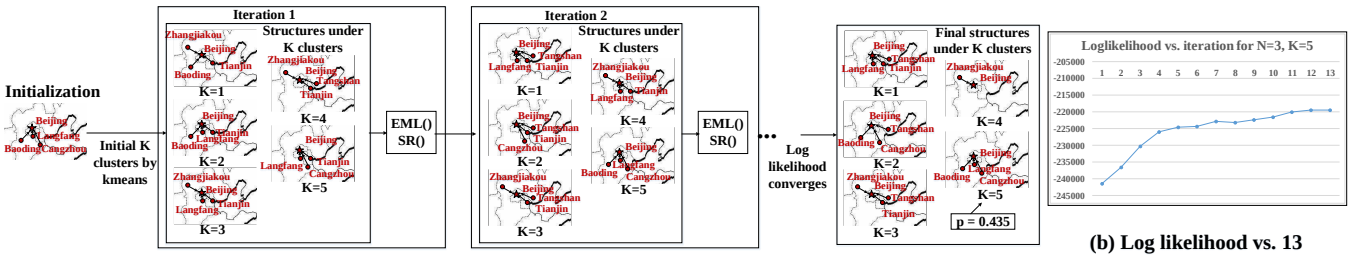

(a) Training process to generate causal pathways under K clusters

(b) Log likelihood vs. 13 iterations of training

Figure 8: An example of learning the causal pathway for Beijing PM2.5, Jan∼Mar under $K = 5$ clusters

We evaluate the empirical performance of our method in this section. All the experiments were conducted on a computer with Intel Core i5 3.3Ghz CPU and 16GB memory. We use MATLAB for our Bayesian learning module, and the open-source MATLAB BNT toolbox [18] for baseline methods.

5.1 Experimental Setup

5.1.1 Data Sets

We use three data sets that contain the records of 6 air pollutants and 5 meteorological measurements:

• North China (NC), with 61 cities, 544 air quality monitoring sensors and 404 meteorological sensors in North China. The latitude and longitude ranges are 34N-43N, 110E-123E.

• Yangtze River Delta (YRD), with 49 cities, 330 air quality monitoring sensors and 48 meteorology sensors. The latitude and longitude ranges are 28N-35N, 115E-123E, respectively.

• Pearl River Delta (PRD), with 18 cities, 124 air quality monitoring sensors and 406 meteorology sensors. The latitude and longitude ranges are 22N-25N, 110E-116E.

The 6 air pollutants are PM2.5, PM10, NO_2, CO, O_3, SO_2, and the 5 meteorological measurements are temperature (T), pressure (P), humidity (H), wind speed (WS), and wind direction (WD). The time span for all data sets is from 06/01/2013 to 19/12/2015. We separate each data set into four groups based on four seasons, and use the last 15 days in each season in year 2014 and 2015 for testing, and the remaining data for model training. We divide each region into 3×3 grids and average the meteorology values within each grid to get the environmental factors $\boldsymbol{E_t}$ for the coupled model.

We conduct experiments at both city level (Section 5.2.2, 5.2.1, 5.2.4) and sensor level (Section 5.2.3). The city-level experiments average value of the sensors in the city to form a pseudo sensor, and discover the pathways among all the cities in three data sets. The sensor-level experiments analyze the causal relationships among sensors in each data set.

5.1.2 Baselines

Since Bayesian-based methods have been well used to learn causal Bayesian structures [11], we choose the most commonly used BN structure learning approaches as baselines to compare with our method. To identify the dependencies among different pollutants, the baselines are deployed to learn the causal structures for each target pollutant.

1. MWST. Maximum Weighted Spanning Tree (MWST) generates an undirected tree structure based on the MWST algorithm [5]. Each time it connects one edge between two nodes with the maximum mutual information. Furthermore, [23] proposed an independency test method to assign a direction to each edge in the tree structure.

2. MCMC. Markov-chain Monte Carlo (MCMC) is a statistical method that also samples from the Directed Acyclic Graph (DAG) space [3]. The method maximizes the score from a set of similar DAGs that add, delete, or reverse connections, and updates the structure in the next iteration.

2. K2+PS. $K2$ is a widely used greedy method for Bayesian structure learning, which selects at most N parents based on the $K2$ score [6] for each variable given the updating order of all the variables. In our case, we use pattern search algorithm [16] to optimize the updating order, thus reducing the search space of casual pathways of different pollutants. Note that the original $K2$ score is defined for discrete variables. Here we use GC_{score} instead for the continuous variables.

5.1.3 Parameter Setting

The parameters of *p-Causality* are set as follows. (1) the support threshold σ; (2) the temporal constraint Δt; (3) the distance threshold δ_g for finding candidate causers; and (4) the correlation threshold δ_p for finding candidate causers; (5) the number of time lags L = 3; (6) and the number of pollutant categories M = 6. When finding causal pathways at city level, we set $\sigma = 0.1$, $\Delta t = 1$ hour, $\delta_g = 200$ km, and $\delta_p = 0.5$. At the station level, all the the parameters are set the same except that $\delta_g = 15$ km to impose a finger granularity for finding candidate causers. K and N are evaluated within the range K $= 3 \sim 7$, and N$= 1 \sim 5$.

5.2 Experimental Results

5.2.1 Inference Accuracy

We first evaluate the effectiveness of our approach via the causal inference accuracy through the causal pathways at city level, which is a 1-hour prediction task based on our proposed GBN-based graphical model. Note this prediction task is not general for all the timestamps, it only predict the future 1-hour based on the extracted pattern-matched periods, indicating the causal inference for the frequent evolving behaviors. Specifically, we first infer the probability $Pr(k)$ of the testing data belong to cluster k. Then, we use the structure and parameters from the trained causal pathways regarding to this cluster to estimate the future pollutant concentration by Eq. 10.

$$P^{est}_{c_m s_0 t} = \Sigma^{\mathrm{K}}_{k=1}(\mu_{0k} + \boldsymbol{PA}(P_{c_m s_0 t})\boldsymbol{A_k})Pr(k) \qquad (10)$$

The accuracy is defined as $\Sigma^{\mathrm{T_{test}}}_{t=1}(P^{est}_{c_m s_0 t} - P^{*}_{c_m s_0 t})/\mathrm{T_{test}}$, where $P^{*}_{c_m s_0 t}$ is the ground truth value and $\mathrm{T_{test}}$ is the number of test cases. Table 2 shows the 1-hour prediction accuracy for our approach *p-Causality*, *p-Causality-n* which is *p-Causality* without integrating the environmental factors, and three baseline methods in three regions of China. The accuracy shown in Table 2 is the average accuracy for four seasons of each data set. The *p-Causality* gets the highest accuracy. The highest inference accuracy for the

three regions are marked with three different colors (orange-> NC, blue->PRD, and green -> YRD) given different parameters $\mathrm{N} = 3, \mathrm{K} = 5$. We note $\mathrm{N} = 3, \mathrm{K} = 5$ provides the best performance for Region NC, while $\mathrm{N} = 2, \mathrm{K} = 3$ generates the best accuracy for YRD and $\mathrm{N} = 2, \mathrm{K} = 5$ for PRD. The optimal number $\mathrm{N} = 2$ or $\mathrm{N} = 3$ for the neighbor sensors also suggest that the air pollution is mainly affected by the most influential sensors in the ST space. We also find the accuracy of Region PRD is 76.7%, lower than Regions NC and YRD (86% and 89.1%). This may be explained by the fact that PRD region in China has relatively low pollution, while our approach targets at capturing the changes of pollutants and identifying their causality, thus may not be quite applicable for prediction tasks of low pollution. We also evaluate the 1-hour prediction accuracy with three well-used time series model, i.e., auto-regression moving average (ARMA) model, linear regression model (LR), and support vector machine for regression with a Gaussian radial basis function (rbf) kernel (represented as SVM-R). *p-Causality* demonstrate more than 10% higher inference accuracy compared with these time series models.

Table 2: Accuracy of PM2.5 1-hour prediction vs. baselines.

Precision (NC; YRD; PRD)	N = 1	N = 2	N = 3	N = 4	N = 5	p-Causality (Optimal K, N)	p-Causality-n	MWST	MCMC	K2+PS	ARMA	LR	SVM-R
K = 3	0.820	0.791	0.824	0.603	0.638	0.860 (K=5, N=3)	0.812 (N=3)	0.691	0.657	0.801	0.785	0.802	0.653
	0.863	0.891	0.850	0.819	0.609								
	0.912	0.863	0.852	0.867	0.821								
K = 4	0.670	0.697	0.732	0.647	0.756								
	0.809	0.852	0.731	0.725	0.712								
	0.606	0.707	0.629	0.596	0.569								
K = 5	0.738	0.815	0.860	0.754	0.588	0.891 (K=3, N=2)	0.793 (N=2)	0.667	0.632	0.751	0.746	0.788	0.735
	0.863	0.837	0.831	0.741	0.755								
	0.634	0.767	0.636	0.599	0.708								
K = 6	0.734	0.627	0.682	0.492	0.613								
	0.878	0.821	0.812	0.708	0.749								
	0.645	0.592	0.595	0.584	0.579								
K = 7	0.754	0.631	0.813	0.722	0.630	0.767 (K=5, N=2)	0.705 (N=2)	0.658	0.58	0.713	0.754	0.794	0.683
	0.846	0.854	0.668	0.678	0.598								
	0.625	0.649	0.578	0.559	0.505								

Table 3: Computation time for training data sets at city-level.

Time (s)	m = 1	m = 2	m = 3	m = 4	m = 5	m = 6
Mine FEPs	2.84	2.70	3.34	3.14	3.47	3.37
Select candidate causers	44.60	61.26	77.42	55.25	63.60	62.50
Generate initial causal pathways	12.48	13.51	14.29	15.32	12.21	13.60
Integrate confounders	--	--	--	--	--	--
Refine causal structures	3954.96	4316.40	3661.68	3987.12	4507.68	4267.92
p-Causality	4014.88	4393.87	3756.73	4060.83	4586.96	4347.39
Only Bayesian learning module	4533.47	4908.83	4286.98	4590.59	5127.16	4874.74
MWST	1572.29	1709.04	1538.4	1591.1	1796.48	1787.25
MCMC	43548.28	47291.59	42053.29	48314.53	50290.17	52014.68
K2 + PS	27290.18	29378.45	26355.63	28213.09	32058.1	33248.55

5.2.2 *Time efficiency*

We also compare the training time of *p-Causality* with baseline methods. Since our approach consists of both pattern mining and Bayesian learning modules, we present the averaged time consumption of training all the three data sets, for each step in the two modules. We also evaluate the overall time consumption of *p-Causality* and *p-Causality* without pattern mining module (with only Bayesian learning). Results show that our approach is very efficient, with the second minimum computation time among all the methods. MWST consumes the minimal time, however, it does not generate satisfactory accuracy for prediction (as in Section 5.2.1). We thus consider our approach provides the best trade-off regarding to accuracy and time efficiency.

5.2.3 *Scalability*

Another superior characteristic of our approach is the scalability. We further identify the causal pathways for air pollutants at sensor level, which is more than ten times as much in the city-level analysis. Our approach provides linear scalability in time with 1.1 hours training time at city level, and 10 hours at sensor level. Note that when the data becomes larger, the pattern mining module can greatly save the computation time for generating the causal pathways. Our approach needs about 10 hours, which is half of the training time of the approach without pattern mining. Moreover, the computation time can be greatly reduced compared to baseline methods. $MCMC$ even cannot compute such large data sets and $K2+PS$ takes up to 14 days, as shown in 9(b). Meanwhile, the accuracy is guaranteed when extending city-level data to sensor-level data, as shown in Figure 9(a).

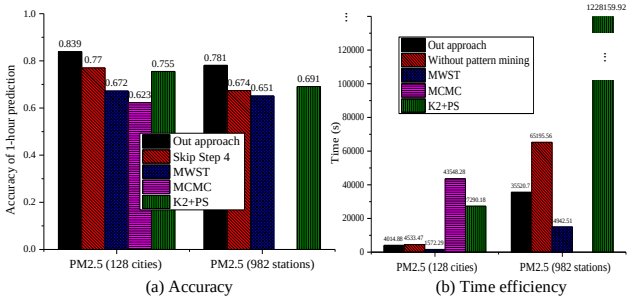

(a) Accuracy (b) Time efficiency

Figure 9: Accuracy and time efficiency at station level

5.2.4 *Case Study*

To analyze the causal pathways for air pollutants, we study two cases corresponding to PM2.5 in specific cities. First we analyze the causal pathways for PM2.5 in the spring of Beijing and in the winter of Shanghai, the period of which are considered as the most heavily polluted season. Then we analyze Beijing PM2.5 before and during the APEC period (1st $\sim$ 14th, Nov, 2014) as a case study for human intervention in causal systems.

1. Beijing and Shanghai. Figure 8 is a real example for the causal pathways for Beijing PM2.5 during Jan$\sim$Mar. We provide the probability for each causal pathway for each cluster, defined as the proportion of labeled timestamps that belong to each cluster. As shown in Figure 8(a), Cluster 5 takes a dominant proportion (43.5%) of time for Beijing PM2.5, indicating the causal pathway during Jan$\sim$Mar most probably come from southern sensors, i.e. Baoding, Cangzhou and Langfang. Actions can be taken to control these pollutants in these cities. We then present the causal pathways for PM2.5 in Shanghai, during Oct$\sim$Dec, which statistically has the highest air pollution concentration. As shown in Figure 10, for PM2.5 in Shanghai, the $\mathrm{N} = 2$ neighborhood cities generally come from the northwest and the southwest. Cluster 2 takes a relative higher proportion (38.4%) of time for Shanghai PM2.5, suggesting the pollutants may be dispersed from PM2.5 in Nantong and SO2 in Zhoushan.

2. Beijing during APEC period. Traditionally, causality is verified via interventions in a causal system. For example, we can verify the effect of a medicine by setting two groups of patients and only giving medicine to the treatment group. However, it is impossible to conduct intervention for air pollutants in the real environment. APEC period is a good opportunity to verify the causality, since the Chinese government shuttered factories in NC, and im-

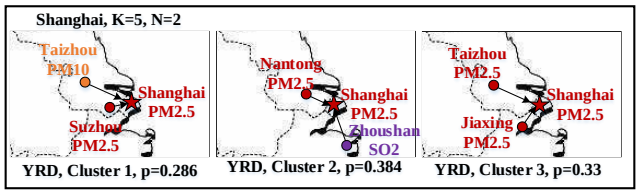

Figure 10: Visualization of final causal pathways for PM2.5 in Shanghai.

plemented traffic bans in and around Beijing [13]. Therefore, we compare the causal pathways for PM2.5 in Beijing before and during the APEC period. To illustrate the propagation of pollutants along the causal pathway, we connect the one-hop pathway to 3-hop as shown in Figure 11(a)(b). The connection originates from the target pollutant, i.e., Beijing PM2.5, and connect its causal pollutants at neighbor cities. Then for each new connected pollutant, we repeat the same procedure for the next hop. The connection stops if in inference accuracy of one target pollutant based on its historical data is higher than based on the historical data of its ST "causers", indicating the pollutant is more likely to be generated locally. Figure 11(a) shows Beijing's PM2.5 is likely to be caused by NO2 in Zhangjiakou, Chengde and Baoding, 2 weeks before APEC period. Further, for example, Baoding's NO2 is mostly influenced by SO2 in Shijiazhuang, Xingtai, and Hengshui. Note that the causal pathways forms "cliques" in the southwestern and northeastern cities to Beijing, which is identical to the locations of the major plants in NC shown in Figure 11(c). However, we notice that the causal pathway cannot be connected into 3-hops during the APEC period, since each "causer" pollutant to Beijing PM2.5 (i.e. NO2 in Chengde and Zhangjiakou, PM10 in Cangzhou) is more likely to be inferred by its own historical data over its ST "causers" in this period. This may suggest the PM2.5 in Beijing during the APEC period are mostly affected by pollutants locally and nearby. The 3-hop causal pathways learnt by three baselines are quite similar, thus we only present the result learned by MWST in Figure 11(d). Our approach has better interpretability. We summarize the discovery for Beijing's air pollution as follows.

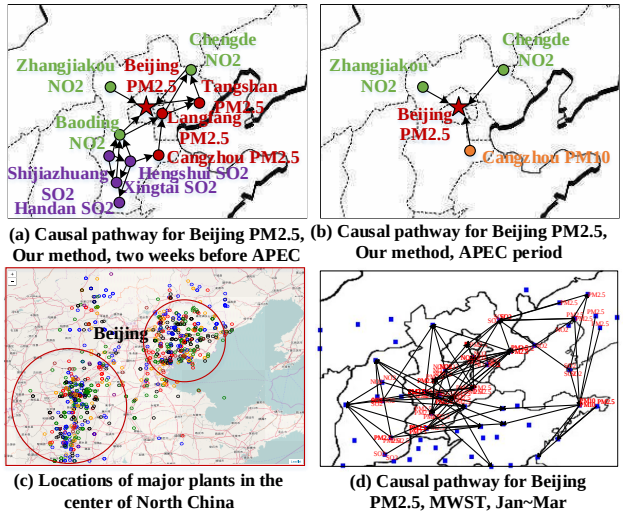

(a) Causal pathway for Beijing PM2.5, Our method, two weeks before APEC

(b) Causal pathway for Beijing PM2.5, Our method, APEC period

(c) Locations of major plants in the center of North China

(d) Causal pathway for Beijing PM2.5, MWST, Jan~Mar

Figure 11: The causal pathways for Beijing PM2.5 before (a) and during APEC period (b), compared with the locations of major plants in Hebei Province, China (c), and the causal pathways learned by baseline method MWST (d).

• Among all the cities within a region, a target pollutant can be mainly affected by only several cities in the ST space. The locations of most influential cities to a target pollutant demonstrate seasonal similarities.

• The causal pathways for PM2.5 in Beijing may come in multiple hops that form "cliques" in certain regions, such as southwest and northeast of Beijing, suggesting superposition or reaction of air pollutants in the corresponding area. While during the APEC period with low pollution level, we did not see multi-hop causal pathways, suggesting the PM2.5 are more likely to be generated locally or nearby within this period.

6. RELATED WORK

In recent years, air pollution analysis has drawn a lot of attention from the data mining community [27]. [28][29] propose data-driven approaches to infer and forecast fine-grained air quality using heterogeneous urban data. [24] estimates the gas consumption and pollutants emission of vehicles, based on the vehicles' GPS trajectories in the road network. Our paper differs from these works in that, we focus on understanding the underlying causal pathways of air pollution, rather than predicting air quality or estimating pollutant emission.

Causal modelling has been an active research direction during the past few years. Existing works on modeling causality for time series data can be classified into two categories. The first category considers a pair of time series, and aims to quantify the strength of causal influence from one time series to another. Researchers have developed different measures for this purpose, such as transfer entropy [2], Granger's causality [10][12]. The second category aim to extract graphical causal relations from multiple time series. [1] combine graphical techniques with the classic Granger causality, and propose a model to infer causality strengths for a large number of time series variables. Pearl's causality model [21] encodes the causal relationships in a DAG [11] for probabilistic inference, with various extensions that incorporating spatial correlation [20] and temporal dependency [19].

Our proposed approach *p-Causality* belongs to the category of using graphical model to detect causalities from multiple time series. It differs from the above works in two aspects: (1) We combine pattern mining and Bayesian learning to make the causality analysis more efficient and robust to the noise present in the input data. (2) Besides the multi-variate time series data, we also consider the impact of confounding given different environmental factors for unbiased causality analysis.

7. CONCLUSION

In this paper, we identified the *ST causal pathways* for air pollutants using large-scale air quality data and meteorological data. We have proposed a novel causal pathway learning approach named *p-Causality* that tightly combines pattern mining and Bayesian learning. Specifically, by extending existing sequential pattern mining techniques, *p-Causality* first extracts a set of FEPs for each sensor, which capture most regularities in the air polluting process, largely suppress data noise and reduce the complexity in the ST space. In the Bayesian learning module, *p-Causality* leverages the pattern-matched data to train a graphical structure, which carefully models multi-faceted causality and environmental factors. We performed extensive experiments on three real-word data sets. Experimental results demonstrate that the causal pathways detected by *p-Causality* are highly interpretable and meaningful. Moreover, it outperforms baseline methods in both efficiency and inference accuracy. For future work, we plan to apply this pattern-aided causal-

ity analysis framework for other tasks in the ST space, such as traffic congestion analysis.

8. REFERENCES

[1] A. Arnold, Y. Liu, and N. Abe. Temporal causal modeling with graphical granger methods. In *KDD*, pages 66–75. ACM, 2007.

[2] L. Barnett, A. B. Barrett, and A. K. Seth. Granger causality and transfer entropy are equivalent for gaussian variables. *Physical Review Letters*, 2009.

[3] J. L. Beck and S.-K. Au. Bayesian updating of structural models and reliability using markov chain monte carlo simulation. *Journal of Engineering Mechanics*, 128(4):380–391, 2002.

[4] D. M. Chickering. Learning bayesian networks is np-complete. In *Learning from data*, pages 121–130. Springer, 1996.

[5] C. Chow and C. Liu. Approximating discrete probability distributions with dependence trees. *Information Theory, IEEE Transactions on*, 1968.

[6] G. F. Cooper and E. Herskovits. A bayesian method for the induction of probabilistic networks from data. *Machine Learning*, 9(4):309–347, 1992.

[7] R. D'Agostino and E. S. Pearson. Tests for departure from normality. empirical results for the distributions of b2 and b1. *Biometrika*, 60(3):613–622, 1973.

[8] I. Flesch and P. J. Lucas. Markov equivalence in bayesian networks. In *Advances in Probabilistic Graphical Models*, pages 3–38. Springer, 2007.

[9] M. A. Gómez, P. M. Villegasa, H. Navarrob, and R. Susia. Dealing with uncertainty in gaussian bayesian networks from a regression perspective. *on Probabilistic Graphical Models*, page 145, 2010.

[10] C. W. Granger. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica: Journal of the Econometric Society*, pages 424–438, 1969.

[11] D. Heckerman, D. Geiger, and D. M. Chickering. Learning bayesian networks: The combination of knowledge and statistical data. *Machine Learning*, 20(3):197–243, 1995.

[12] K. Hlaváčková-Schindler, M. Paluš, M. Vejmelka, and J. Bhattacharya. Causality detection based on information-theoretic approaches in time series analysis. *Physics Reports*, 441(1):1–46, 2007.

[13] K. Huang, X. Zhang, and Y. Lin. The "apec blue" phenomenon: Regional emission control effects observed from space. *Atmospheric Research*, 164:65–75, 2015.

[14] A. Keats, E. Yee, and F.-S. Lien. Bayesian inference for source determination with applications to a complex urban environment. *Atmospheric environment*, 41(3):465–479, 2007.

[15] S. Lee, W. Liu, Y. Wang, A. G. Russell, and E. S. Edgerton. Source apportionment of pm 2.5: Comparing pmf and cmb results for four ambient monitoring sites in the southeastern united states. *Atmospheric Environment*, 42(18):4126–4137, 2008.

[16] R. M. Lewis and V. Torczon. A globally convergent augmented lagrangian pattern search algorithm for optimization with general constraints and simple bounds. *SIAM Journal on Optimization*, 12(4):1075–1089, 2002.

[17] J. Lin, E. Keogh, S. Lonardi, and B. Chiu. A symbolic representation of time series, with implications for streaming algorithms. In *SIGMOD*, 2003.

[18] K. Murphy et al. The bayes net toolbox for matlab. *Computing science and statistics*, 33(2):1024–1034, 2001.

[19] K. P. Murphy. Dynamic bayesian networks. *Probabilistic Graphical Models, M. Jordan*, 7, 2002.

[20] H. Nguyen, W. Liu, and F. Chen. Discovering congestion propagation patterns in spatio-temporal traffic data. In *the 4th International Workshop on Urban Computing (UrbComp)*, 2015.

[21] J. Pearl. Causality: models, reasoning and inference. *Economet. Theor*, 19:675–685, 2003.

[22] J. Pei, J. Han, B. Mortazavi-Asl, H. Pinto, Q. Chen, U. Dayal, and M. Hsu. Prefixspan: Mining sequential patterns by prefix-projected growth. In *ICDE*, pages 215–224, 2001.

[23] G. Rebane and J. Pearl. The recovery of causal poly-trees from statistical data. *arXiv preprint arXiv:1304.2736*, 2013.

[24] J. Shang, Y. Zheng, W. Tong, E. Chang, and Y. Yu. Inferring gas consumption and pollution emission of vehicles throughout a city. In *KDD*, 2014.

[25] W. Sun, P. Wang, D. Yin, J. Yang, and Y. Chang. Causal inference via sparse additive models with application to online advertising. In *AAAI*, pages 297–303, 2015.

[26] C. Zhang, Y. Zheng, X. Ma, and J. Han. Assembler: Efficient discovery of spatial co-evolving patterns in massive geo-sensory data. In *KDD*, pages 1415–1424. ACM, 2015.

[27] Y. Zheng, L. Capra, O. Wolfson, and H. Yang. Urban computing: concepts, methodologies, and applications. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 5(3):38, 2014.

[28] Y. Zheng, F. Liu, and H.-P. Hsieh. U-air: When urban air quality inference meets big data. In *KDD*, pages 1436–1444. ACM, 2013.

[29] Y. Zheng, X. Yi, M. Li, R. Li, Z. Shan, E. Chang, and T. Li. Forecasting fine-grained air quality based on big data. In *KDD*, 2015.

[30] J. Y. Zhu, C. Sun, and V. O. Li. Granger-causality-based air quality estimation with spatio-temporal (st) heterogeneous big data. In *Computer Communications Workshops (INFOCOM WKSHPS), 2015 IEEE Conference on*, pages 612–617. IEEE, 2015.

[31] J. Y. Zhu, Y. Zheng, X. Yi, and V. O. Li. A gaussian bayesian model to identify spatio-temporal causalities for air pollution based on urban big data. In *Computer Communications Workshops (INFOCOM WKSHPS), 2016 IEEE Conference on*. IEEE, 2016.